



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

I also certify that the attached copy of the request for grant of a Patent (Form 1/77) bears an amendment, effected by this office, following a request by the applicant and agreed to by the Comptroller-General.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 2 December 2003

Patents Form 1/77

Patents Act 1977
(Rule 16)

The
Patent
Office

1/77

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference	SIBU / P21280GB		
2. Patent application number (The Patent Office will fill in this part)	9918611.6		
3. Full name, address and postcode of the or of each applicant (underline all surnames)	Sibelius Software Ltd. 75 Burleigh Street Cambridge CB1 1DJ United Kingdom	66-68 HILLS ROAD CAMBRIDGE CB2 1LA UNITED KINGDOM	AG A/L 21/7/00
Patents ADP number (if you know it)	07716970002		
If the applicant is a corporate body, give the country/state of its incorporation	United Kingdom	7716970001	
4. Title of the invention	MUSIC DATABASE SEARCHING		
5. Name of your agent (if you have one)	ERIC POTTER CLARKSON PARK VIEW HOUSE 58 THE ROPEWALK NOTTINGHAM NG1 5DD		
"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)			
Patents ADP number (if you know it)	1305010		
6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number	Country	Priority application number (if you know it)	Date of filing (day / month / year)
7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	Number of earlier application		Date of filing (day / month / year)
8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:	YES		
a) any applicant named in part 3 is not an inventor; or			
b) there is an inventor who is not named as an applicant, or			
c) any named applicant is a corporate body.			
See note (d))			

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description 21
Claims(s) 4
Abstract 1
Drawing(s) 2 + 2

10. If you are also filing in any of the following, state how many against each item.

Priority Documents 0
Translations of priority documents 0
Statement of inventorship and right to grant of a patent (Patents Form 7/77) YES
Request for preliminary examination and search (Patents Form 9/77) NO
Request for substantive examination (Patents Form 10/77) NO
Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application.

Signature *Eric Potter Clarkson*
ERIC POTTER CLARKSON

Date
6 August 1999

12. Name and daytime telephone number of person to contact in the United Kingdom 0115 9552211

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 01645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

DUPLICATE

MUSIC DATABASE SEARCHING

The present invention relates to search engines and databases, and in particular to search engines adapted to search for particular musical sequences or phrases in a database of recorded or encoded sound files in a computer system.

The invention relates to searching of databases of varying types. The database could be restricted in size, scope and file format such as a publisher's compact disc catalogue. Alternatively, the database search might be extensive in size and scope, may be widely distributed on a network, and may incorporate many different file types. One example would be an internet search.

In many circumstances, it is desirable to be able to search a music database for a specific piece of music, based solely upon knowledge of a portion of a tune or musical sequence from a piece of music. Otherwise, more detailed conventional bibliographical information such as the title of the work, composer, publisher, lyrics etc. must be provided to effect a search, and such details might not always be known to the searcher.

Known problems associated with searching for a piece of music in a database of music files, based only on knowledge of a tune, are many and varied.

Firstly, a suitable input device for providing the known tune as search criteria to a computer system is required. This is difficult, since the input of music to a computer system, via conventional computer input devices (such as keyboards) by unskilled users, is not straightforward.

Secondly, a method of comparing the search criteria with the complex patterns likely to be found in a computer-based music file is difficult because the precise location of the recognisable tune within the complexities of recorded or encoded sound is not known. A variety of file
5 types such as MIDI files, MP3 files, WAV files, sequencer files, scorewriter files or files in other suitable formats must be accommodated.

Throughout the present specification, the expression "music file" will be used to encompass all forms of electronically, magnetically or optically
10 stored computer-readable files which contain a digital representation of music, from which musical pitch data can be extracted. These representations could relate to encoded recorded sound such as in an MP3 file format, or to coded instructions for creating sound such as a MIDI file format.

15 Throughout the present specification, the expression "tune" will be used to indicate a sequence of note pitches, preferably single notes rather than chords, which can form the basis of search criteria. Throughout the present specification, the expression "melody" will generally be used to
20 refer to sequences of note pitches in portions of a music file to be searched which are likely to be locations where a search tune will be found, eg. vocal lines, or solo instrumental lines.

In the prior art, it has been suggested that search criteria can be specified
25 by a relatively simple method of providing a sequence of musical contours. These musical contours describe relative pitch transitions and simply indicate whether each successive note is higher, lower or the same as the preceding note. This format lends itself to easy keyboard input by a user simply providing a character string such as "DUDRUDDUUDUUDUDR"
30 where "D" represents a downward transition, "U" indicates an upward

transition and "R" indicates a repetition of the previous note pitch. Such techniques have found some success with specially prepared databases but are limited by their inaccuracy and input of search criteria is still somewhat awkward for the unskilled user. In addition, such techniques are not particularly suited to searching general music files.

It is an object of the present invention to provide a method and apparatus for providing musical search criteria as input to a search engine, in a manner which is easy to use by the unskilled or non-expert user.

It is a further object of the present invention to provide a method and apparatus for applying musical search criteria to a database to obtain a match against target music files in a computer storage medium.

It is a further object of the present invention to provide a method and apparatus for structuring music files in a computer system database in order to enable rapid or efficient searching thereof for specified search criteria comprising a tune.

According to one aspect, the present invention provides an apparatus for effecting a search through a database of music files, comprising:

input means, for providing as input search criteria comprising a tune as a sequence of melodic intervals;

comparing means, for comparing said sequence of melodic intervals with selected portions of a plurality of computer-readable music files; and

output means, for providing as output a list of possible matches of said search criteria with ones of said plurality of computer-readable music files.

According to another aspect, the present invention provides an apparatus for indexing a music database comprising:

means for identifying relevant selected portions of a plurality of computer-readable music files by applying selection criteria to identify
5 portions of the files likely to contain tunes; and

means for tagging said music files to identify positions corresponding to said relevant selected portions.

According to another aspect, the present invention provides an apparatus
10 for indexing a music database comprising:

means for identifying relevant selected portions of a plurality of computer-readable music files by applying selection criteria to identify portions of the files likely to contain tunes; and

means for generating an index of said music files containing
15 information representative of said relevant selected portions.

According to another aspect, the present invention provides a method for effecting a search through a database of music files, comprising:

providing as input, search criteria comprising a tune as a sequence of
20 melodic intervals;

comparing said sequence of melodic intervals with selected portions of a plurality of computer-readable music files; and

providing as output a list of possible matches of said search criteria with ones of said plurality of computer-readable music files.

25

Embodiments of the present invention will now be described by way of example and with reference to the accompanying drawings in which:

Figure 1 shows a schematic diagram of apparatus for conducting a music search in a database, according to an embodiment of the present
30 invention; and

Figure 2 shows a flow diagram of a music search method according to one aspect of the present invention.

5 With reference to figure 1 there is shown a computer system 1 suitable for implementing the music search method of the present invention. Preferably, the computer system comprises conventional hardware in the form of a processor 2, memory 3, monitor 4, keyboard 5 and microphone 6. Preferably, the computer system 1 also includes a MIDI compatible music keyboard 7 and appropriate modem links 8 to external data databases
10 9.

With reference also to figure 2, the search procedure of one embodiment of the invention will now be described in connection with further details of the computer system 1 of figure 1.

15

A first step 20 is to input search criteria relating to the tune which is being sought. Preferably, this is effected by a user singing, humming or whistling a tune to the computer using microphone 6. In step 21, this audio input is used to generate an audio file. In step 22, a pitch recognition
20 engine 11 in processor 2 is then used to analyse the audio file and identify the pitch, and preferably also the duration, of successive notes in the tune.

In an alternative embodiment, it will be understood that input of the tune can also be readily achieved using a number of other methods in place of
25 step 21, and usually also step 22. For example, the MIDI keyboard 7 can be used to play the tune directly into a MIDI file. The conventional computer keyboard 5 or a mouse can be used to type in note names (eg. C, E, F etc). A tune may also be selected from a portion of an existing stored music file, eg. MIDI file. This technique can be useful if a large number
30 of searches are to be made, and could include automatic selection of many

parts of a number of music files, if comparison between many files in a database is desired.

Once the succession of note pitches defining the tune to be used as search
5 criteria has been entered into the computer, the next step (figure 2, step 23)
is to determine a sequence of melodic intervals from the note pitches. A
query definition procedure 12 uses a melodic interval generator 13 to
specify the sequence of melodic intervals used in the query.

10 A melodic interval is defined as the pitch interval between a note and a
preceding note. Preferably, the preceding note is the immediately
preceding note although the melodic interval could be defined relative to
the first note of a tune or segment of a tune. This latter option is less
resilient to the pitch drifting sharp or flat where the tune has been input
15 using audio input, eg. by singing, humming or whistling.

Preferably, for audio input, the melodic intervals are quantized to reduce
pitch errors, analogous to "snap-fitting" graphics lines or points to a grid.
(This may also be appropriate for inaccurate MIDI input, eg. where a user
20 plays a slightly wrong note.) The optimum degree of quantization is best
found statistically from a representative set of music files and search tunes,
but other quantization strategies used in the present invention include any of
the following.

25 Quantization to half-steps (semitones) is preferably used if the database is
reliable and the search tune is specified by a trained musician using MIDI
keyboard or other reliable input device.

Quantization to whole-steps (tones) or minor thirds if the database or search
30 tune is less reliable, e.g. sung or hummed by an untrained user.

Quantizing to a scale, e.g. diatonic, can be used, if this can be reliably specified in, or deduced from, the music files and search tune.

- 5 In a preferred embodiment, not only are melodic intervals used in the search criteria but also, in step 24, the query definition procedure 12 uses a rhythmic interval generator 14 to determine a sequence of rhythmic intervals from the notes in the tune.
- 10 Rhythmic intervals are defined as the duration of a note divided by the duration of the preceding note. So a rhythmic interval of 2 denotes a pair of notes of which the second is twice as long as the first. As with melodic intervals, a note's rhythmic interval could be relative to the first note of the tune or to the first note of a segment of a tune, though this would be less
15 resilient to acceleration or deceleration within a tune.

Rhythmic intervals are best used when the search tune is input in a way which accurately records rhythm (eg. singing or MIDI keyboard). Rhythmic intervals are preferably used as second order search criteria,
20 carrying less weight than the melodic intervals, because there is a tendency for them to be much less accurate. This may be because the search tune has only approximate rhythm as input, or because there are often rhythmic variations between different arrangements of music.

- 25 Preferably, rhythmic intervals should be coarsely quantized into just a few ranges. These ranges should be statistically chosen such that rhythmic intervals rarely fall near a range boundary. For instance, quantizing rhythmic intervals into 0-1, 1-2 and 2+ would be very poor as many rhythmic intervals will be close to 1 and would be allocated between the
30 first two ranges almost at random.

In a preferred embodiment, as shown in step 25, further search criteria can also be specified, in addition to melodic intervals and rhythmic intervals, to further refine a search. For example, the query definition procedure may
5 facilitate input of any or all of the following general bibliographic or text information search criteria if such information will ordinarily be found in, or linked to, the music files to be searched: a) lyrics; (b) title and composer; (c) artist; (d) filename; (e) date of composition – the user may well know this approximately (e.g. decade for pop/rock songs, century for
10 classical music) without knowing the title or composer.

Once all of the search criteria have been specified in the query definition procedure 12 (step 26), a comparison procedure 15 is initiated.

15 Relevant features in the search criteria are compared with relevant features in each music file in a database 9 or 10. File indexes may be used if available, to be discussed hereinafter.

Preferably, melodic interval sequences and/or rhythmic interval sequences
20 are segmented (step 27) during the comparison process (step 29) by a segmentation algorithm 16. The main purpose of segmentation is to provide resilience against there being an extra note in the search tune which is not in the target file, or vice versa. If this discrepancy occurs early on in the search tune, then comparing it on a note-by-note basis with the target
25 file will produce poor matches for all subsequent notes.

However, if the search tune and target music file are divided into segments of a few notes in such a way that an error in a segment will not affect later segments, then a note omitted or added will only affect the score of one
30 segment and will not seriously affect the overall match.

Segments should be a few notes long, preferably 3 to 7 notes. If segments are too short then they will not be distinctive of a particular tune, and false positives will occur by virtue of a different tune containing the same segments (unless a higher score is given for segment order). If segments are too long, then errors in a search tune segment will produce a low score (unless near-matches are scored higher than poor matches).

The ideal segment length and segmenting algorithm can be derived statistically from a representative sample of music files and search tunes.

According to preferred embodiment, segmentation algorithms may include any of the following techniques.

1. Segmentation into variable-length segments based on local context within the tune. For example, a segment boundary could occur at each local maximum or minimum pitch or rhythmic interval. It is important that segmentation depends only on local context, otherwise a note error in the search tune could affect multiple segments and produce too low a score. The algorithm must also avoid producing excessively long or short segments for non-standard kinds of tune.

2. Segmentation into bars or groups of bars, if these can be reliably identified both in the music files and in the search tune and are likely to be the same in both. An example of this is if the music files and search tune are both in the form of music notation, such as with scorewriter files.

3. Segmentation into overlapping, fixed-length segments, typically of 3-7 notes. For example, the sequence of notes A B C D E F could be segmented into the following segments: ABC, BCD, CDE, DEF. This is

similar to segmenting into bars but does not require the bar length or barline position to be identified. A potential disadvantage of this method is that it produces many segments and so is relatively inefficient.

- 5 It is possible to produce the same effect as segmentation in another way, but this is relatively inefficient. It would require (i) comparing the search tune with the music file note-by-note, starting at all possible start-points within the music file (if the start of the tune has not been reliably detected in the music file), and (ii) shifting the comparison point forward or
10 backward one or more notes whenever a note-match failed in order to try to get the search tune and music file back in step.

It is noted that the ideal segment length and the best segmenting algorithm may vary according to the type of music being searched. Thus in another
15 embodiment, the query specifies the segmenting algorithm to be used. This may be done directly by the user, or, more preferably, indirectly by the user indicating the type of music being searched for, eg. rock, jazz, country, classical, etc.

- 20 In step 28, the comparison procedure systematically retrieves files from the databases 9, 10 and preferably performs the comparison operation (step 29) on the basis of segments of the search criteria tune and segments of the file. The comparison operation may also include other search criteria, such as text, as discussed earlier.

25 Each music file compared is given a score which is higher for a better match (step 30). Scores may distinguish the closeness of match between segments, or may just assign a single high score (e.g. 1) for an exact match and a single low score (e.g. 0) for an inexact match or non-match.

30

Statistics may be used to model likely errors and score inexact matches accordingly. For example, the interval of a tritone is hard to sing correctly, so if it occurs in a music file but does not match the search tune then this may be penalised less than easier intervals which fail to match.

5

A higher score may also be assigned for an exact or close match between the order of segments (to rule out melodies which have similar groups of notes but in a different order).

- 10 A higher score may be assigned if the search tune is found to start at (or near) the start of a relevant selected portion in the music file, as it is likely the search tune is the start of a melody in the search file, rather than the middle or end.

- 15 Scores for separate features of a music file may be combined by adding or otherwise.

- To increase the speed of searching large databases, there may be one or more coarse comparisons in which some reliable features are compared in
20 order to exclude most of the music files, followed by a more detailed comparison of all features in order to produce a reliable score for each file not already excluded.

- In order to facilitate searching of varying music file types some pre-
25 processing of the music files being searched may be necessary, or may simply be preferable to speed up the matching procedure.

- If the music files are audio-based, ie. comprising a representation of recorded sound, a pitch recognition engine similar to that described in
30 connection with the query definition could be used to first determine the

notes of possible melodies contained in the music file. These melodies are then compared with the tune of the search criteria, as discussed above.

5 A music file will typically contain not only melody but also other notes playing simultaneously forming an accompaniment. The melody will typically consist of a tune, typically repeated for two or more verses, and often preceded and/or followed by music other than the tune, e.g. an introduction or coda.

10 To reduce the amount of music to be searched, it is desirable to select portions of the music file which are more likely to include the tune being searched for, and thereby exclude accompanying instruments and sounds other than the melodies which may contain the tune.

15 In some music file formats, e.g. MIDI files, sequencer files and scorewriter files, the file is typically divided into separate streams of pitches called tracks, channels or staves. One track, channel or staff typically contains the music for one instrument. The melody is often on the same track, channel or staff throughout the file, or only changes track, channel or staff
20 infrequently within the file.

In other file formats, e.g. WAV files, the file is typically not so divided. However, algorithms can be used to separate out streams of pitches which probably correspond to individual instruments. Once this has been done,
25 streams which may contain the tune can be identified in similar ways as for tracks, channels or staves.

For audio files, this separation of streams can be achieved by performing a Fourier analysis to turn the waveform into frequencies, then, using
30 Bayesian or other techniques to hypothesise which of these frequencies are

genuine notes (typically fundamental harmonics of notes) and which are just higher harmonics of such notes. Durations can then be derived for the notes. Successive notes which are of similar timbre and/or pitch and/or which are connected by smooth transitions can then be regarded as
5 continuous streams of pitches on the same instrument. We will refer to these as 'streams' below.

The following criteria can be used for identifying the track, channel, staff or stream containing the melody, or for tracking the melody if it switches
10 between different tracks, channels or staves:

- a) If one track/channel/staff/stream has lyrics, it is very likely to be the melody
- b) The melody is usually on track 1 in type 1 MIDI files, or the top staff in scorewriter files
- 15 c) In MIDI files, the melody is often on channel 1, or if this is silent then on channel 2
- d) In MIDI files, the melody is very unlikely to be on channel 10 (unpitched percussion)
- e) The melody is usually the highest pitch most of the time, particularly
20 if the highest pitch is on the same channel, track or staff continuously for long stretches of music
- f) The melody is very unlikely to contain more than a few melodic intervals greater than an octave
- g) If a track, channel, staff or stream is not almost entirely diatonic, it
25 is probably not the melody
- h) The track, channel or staff containing the melody probably consists of single notes rather than chords
- i) The melody probably has an average note length of between, say, 0.15 and 0.8 seconds (i.e. is probably not extremely fast or slow)

- j) The highest note in the melody is probably B above middle C or higher, and the lowest note is probably D a ninth above middle C or lower
 - k) If the track or staff name or first text item contains 'lead', 'melody',
5 'vox' or 'vocals', then it is very likely to be the melody
 - l) If the track or staff name contains 'bass' then it is probably not the melody
 - m) If a sound used in much or all of the track/channel/staff/stream is strange then it is probably not the melody (e.g. for MIDI files, if the
10 first program change is General MIDI pizzicato, gunshot, timpani or bass guitar, it is probably not the melody)
 - n) If a channel, track or staff is playing for less than a predetermined percentage of the file's duration, it is probably not the melody. In a preferred embodiment, the predetermined percentage is
15 approximately 70%
 - o) The melody is probably fairly continuous for several seconds at a time (e.g. it is extremely unlikely to consist of rests punctuated by an occasional isolated note)
 - p) The melody is probably at least as loud as all other instruments
20 (except unpitched percussion)
 - q) The melody is unlikely to consist of a short melodic or rhythmic segment exactly repeated many times in succession (such as may occur in accompaniments).
- 25 Thus, in a preferred embodiment, a more sophisticated algorithm is possible to first identify relevant portions of the music files which may contain the tune being sought, prior to the comparison step 29 taking place. Preferably, in step 32, a set of selection criteria, exemplified in items a to q above are preferably used to identify relevant selected portions of the music
30 file being searched which are melodies likely to contain search tunes.

The relative scores or weights to allocate to these selection criteria are best determined empirically (e.g. by Bayesian statistics) from a large set of music files representative of those likely to be searched. A score
5 combining most or all of the above criteria would be highly reliable even though the individual criteria are not. The selection criteria may be selected by the user according to the type of music known to be in the database, as discussed in relation to the segmentation process.

10 If no one portion of the file scores significantly higher than others according to these criteria, several melodies in the file may need to be identified as relevant selected portions of the file and searched.

If several melodies in the file are identified, probabilities or other weights
15 based on the above criteria indicating how likely each melody is to be the actual search tune could also be stored, and used when scoring matches against the search tune. Scoring criteria could include a weighting factor giving an indication as to how popular the matched melody is. For example, if the melody is a recent pop song, it would be afforded a higher
20 probability of a true match than a piece by an obscure composer. Popularity could be automatically determined on the basis of frequency of hits on the database, eg. number of file accesses or search matches.

Additionally, it may be desirable to identify the point at which each melody
25 (or relevant selected portion) starts, as this is most likely to match the search tune. Criteria for selection of portions of a music file corresponding to the start of melodies which are likely to include a search tune are:

- i) If there are lyrics, the start of the lyrics is very likely to be the start
30 of a tune.

- ii) The start of the first instance of a passage of reasonable length which is later repeated (i.e. the first of two or more verses) is likely to be the tune.
- iii) Failing these, the start of the music in the channel, track, staff or stream which contains a melody is quite likely to be the start of the tune.

It will be understood that the use of the selection criteria indicated above, to identify likely relevant selected portions of music files to be compared against the search criteria, could be done in real time as the search is progressing. However, more preferably, the application of the selection criteria is carried out on the database files prior to searching. A selection procedure provides this function, either in real time during searching, or independently prior to searching.

15

It will be understood that the relevant selected portions identified by the selection criteria can also be then used as tunes for search criteria, if it is desired, for example, to search a database for any two similar tunes. Such would be useful for looking for potential cases of copyright infringement.

20

In a preferred embodiment, the database music files are examined by an indexing / tagging procedure 17 to identify relevant selected portions, according to the selection criteria a) - q) and i) - iii) above, before any searching takes place. This information is then provided as an index to the file or as a set of tags to identify positions in the file corresponding to relevant selected portions. The index information may be stored separately from or linked to the respective music file, or as header information thereto, for example.

This indexing, or tagging, of the files is particularly important in very large databases in order to speed up searching of the database.

5 The indexing process discards irrelevant features (ie. parts which are clearly not melodies) from the music files and thereby reduces the amount of material to be searched. The index entry for a file is hence a much-abbreviated form of the original file.

10 The first stage of indexing is to identify melodies (ie. candidate tunes) and eliminate any music which is definitely not a melody, e.g. the accompaniment or introduction.

The second stage is to segment the melodies into short groups of notes.

15 The third stage is to extract relevant features from the melodies, primarily melodic intervals. Features other than the melodies can also be indexed to aid matching. For example, these additional features may include: a) lyrics (e.g. in a MIDI file), especially the lyrics at the start of the possible tunes; (b) title and composer (e.g. in a MIDI file); (c) artist; (d)
20 filename, which is often an abbreviation of the music's title, e.g. YELLOW for "Yellow Submarine"; (e) date of composition.

Indexing is preferably carried out automatically but can be partly manual for some applications. For instance, if creating a database of music clips
25 from CDs so that consumers can find CDs in a store, it may be desirable for a human to create a music file corresponding to each CD containing just the first few bars of the melodies of the most important tracks on the CD. These music files form the database which is searched.

Alternatively, tagging files may comprise marking relevant features within an existing file, such as tagging the start of the melodies in a MIDI file. Tagging may also include specifying the time signature and barline positions to assist segmentation during the comparison procedure.

5

Index files can be in any suitable format. However, if an index is in a suitable text format then matching can be done by a normal text search engine, provided the search criteria are encoded similarly.

- 10 In the case of recorded sound files which are being searched, it will be understood that the process for identifying relevant selected portions of music files, either for searching, indexing or tagging, may also need to be preceded by quantization processes as discussed earlier in connection with the determination of the tune to be used as search criteria.

15

- A single web page or text file can be generated as the index entry for a music file. Melodic intervals and rhythmic intervals can be encoded as letters and divided into space-separated 'words' representing segments. For Internet-wide searches, it may be desirable to design the encoding such that
- 20 the 'words' formed are unlikely to occur in normal text and produce false matches in normal text pages; e.g. vowels should be excluded from the encoding.

- If correct segment order is to be scored more highly than incorrect order,
- 25 this can be represented by a text phrase search (which requires words to be in the specified order) or a 'nearness' search. Lyrics, title, artist etc. can be included as ordinary text.

- A number of other strategies for improving the accuracy of the comparison
- 30 procedure may be included.

Repeated or tied notes: the search tune and the target music file may differ because the arrangements or lyrics have tied notes in one and repeated notes in the other. It may be also hard to detect repeated notes in audio,
5 especially if there is no intervening consonant. Differences between repeated and tied notes can be eliminated by treating all repeated notes as tied.

Chords: though tracks, channels or staves containing possible tunes
10 probably contain no or few chords, any chords which do occur could be replaced by the top note of the chord.

Rests in mid-tune: these may vary between arrangements (e.g. due to staccato or legato). Rests can be eliminated by regarding a note followed
15 by a rest as ending at the start of the next note (if the next note starts within, say, 1 second).

Pitch bending (e.g. in MIDI files): this can be ignored, but will typically be eliminated by quantization of melodic intervals as discussed earlier.

20

Octave errors in audio: the difficulty in identifying the octave when performing pitch recognition on audio data can cause notes to be regarded as one or more octaves out. This can be eliminated by transposing all notes to within a single octave range.

25

The computer system 1 may generate the output results from the search in a number of ways. Preferably, in step 30, each music file is awarded a score based on nearness of match of the search criteria with the relevant selected portions of the file. The user will then be presented with a ranked list of
30 the highest scoring files (step 33) together with information about the files

from the database (eg. title, artist etc.). With tagging and/or indexing of files, it is also readily possible to include, or link to, a small clip of the music file to play in order that the user can readily verify whether the melody located therein corresponds to the tune in the search query.

5

There are a large number of applications of the present invention, such as: (a) finding music files on the Internet; (b) finding CDs in music shops and Internet sites; (c) finding music in databases used by rights collecting agencies, advertising agencies, libraries, film/TV companies, etc; (d) 10 comparing music files within a database against each other to find duplicates, different arrangements, or copyright infringement.

Two exemplary applications are given below.

- 15 1. In an Internet search engine for MIDI files, a computer has indexed all MIDI files on the Internet using a web crawler. The user plays a search tune on a music keyboard shown on the screen using a mouse. The computer plays the search tune back before searching, to confirm that it has been input as intended. The user also types in a few words of lyrics from 20 the tune, then clicks a 'Search' button. The computer performs a coarse comparison between the search criteria and its index to identify around 500 files which may contain the search tune. The computer performs a finer comparison to assign scores to these files. The computer lists the highest-scoring 20 files by filename. The user can click on any of these filenames 25 to hear the MIDI file and check whether it contains the tune in question. The computer plays from the point in each file where the index or tag indicates the melody or candidate tune begins. The user can then click on a link to jump to the web page which contains the target MIDI file, so that the user can download it.

30

2. In a search engine for CDs in a music store, for each of the best-selling CDs in a music store, a MIDI file has been generated containing the first few bars of each tune on the CD, and the CD's tracks have been converted into separate MP3 files. A consumer sings a search tune into a microphone on an in-store kiosk. The computer in the kiosk converts this audio data into a stream of note pitches and note durations. The computer matches this against its indexed database of MIDI files. For the highest-scoring file, the computer lists the matching CD's title, artist, price and location, and starts playing the matching track's music from the separate MP3 file. The consumer can press a button to skip to the next-highest scoring match.

CLAIMS

1. Apparatus for effecting a search through a database of music files,
comprising:
 - 5 input means, for providing as input search criteria comprising a tune
as a sequence of melodic intervals;
comparing means, for comparing said sequence of melodic intervals
with selected portions of a plurality of computer-readable music files; and
output means, for providing as output a list of possible matches of
10 said search criteria with ones of said plurality of computer-readable music
files.
2. Apparatus according to claim 1 wherein said input means comprises
a microphone into which a user can sing, hum or whistle said tune.
- 15 3. Apparatus according to claim 1 wherein said input means comprises
a MIDI keyboard for playing the tune.
4. Apparatus according to claim 1, claim 2 or claim 3 wherein the input
20 means further includes pitch recognition means for identifying each melodic
interval between a succession of musical pitches input as said tune.
5. Apparatus according to claim 4 wherein said pitch recognition means
further includes quantization means for determining a closest chromatic
25 interval, a closest whole tone interval, or a closest minor third interval
between two successive musical pitches.
6. Apparatus according to claim 4 wherein said pitch recognition means
further includes quantization means for determining a closest major, minor
30 or other scale to which successive musical pitches will fit.

7. Apparatus according to any preceding claim further including means for determining, from said input sequence of melodic intervals, a succession of rhythmic intervals and using said succession of rhythmic intervals as further search criteria.
8. Apparatus according to any preceding claim further including means for providing as input further search criteria comprising text information.
9. Apparatus according to any preceding claim wherein said comparing means includes means for comparing a plurality of segments of said tune with said selected portions of said plurality of computer-readable music files, and wherein said output means bases the likelihood of a match based on the number of separate segments for which a possible match is indicated.
10. Apparatus according to claim 9 wherein said segments of the search tune are defined as overlapping note sequences.
11. Apparatus according to any preceding claim wherein said comparing means includes means to identify relevant selected portions of a plurality of computer-readable music files by applying selection criteria to identify portions of the files likely to contain tunes.
12. Apparatus according to claim 11 wherein said relevant selected portions of said music files are stored in an index.
13. Apparatus according to claim 12 wherein said relevant selected portions stored in said index are encoded as text, said input means further

including means for encoding said sequence of melodic intervals as a text string, said comparing means comprising a text search engine.

14. Apparatus according to claim 11 wherein the location, in said
5 computer-readable music files, of said relevant selected portions of said music files are indicated by one or more tags, said comparing means adapted to locate said tags.

15. Apparatus for indexing a music database comprising:
10 means for identifying relevant selected portions of a plurality of computer-readable music files by applying selection criteria to identify portions of the files likely to contain tunes; and
means for tagging said music files to identify positions corresponding to said relevant selected portions.

15
16. Apparatus for indexing a music database comprising:
means for identifying relevant selected portions of a plurality of computer-readable music files by applying selection criteria to identify portions of the files likely to contain tunes; and
20 means for generating an index of said music files containing information representative of said relevant selected portions.

17. A method for effecting a search through a database of music files, comprising:
25 providing as input search criteria comprising a tune as a sequence of melodic intervals;
comparing said sequence of melodic intervals with selected portions of a plurality of computer-readable music files; and
providing as output a list of possible matches of said search criteria
30 with ones of said plurality of computer-readable music files.

18. A computer program product, comprising a computer readable medium having thereon computer program code means adapted, when said program is loaded onto a computer, to make the computer execute the
5 procedure of claim 17.

19. Apparatus substantially as described herein with reference to the accompanying drawings.

ABSTRACT

MUSIC DATABASE SEARCHING

- 5 A computer-based system for effecting a search through a database of music files includes an input means, for providing as input search criteria comprising a tune as a sequence of melodic intervals. The input may be by microphone or MIDI keyboard. The search criteria are compared with sequences of melodic intervals from selected portions of a plurality of
- 10 computer-readable music files. The music files may have been indexed or tagged to identify relevant selected portions which are likely to correspond to melodies therein. An output is generated comprising a list of possible matches of said search criteria with ones of said plurality of computer-readable music files, preferably in ranking order.

Figure 1




FIG. 1

FIG 2

